5. CALCULATING ITEM CALIBRATIONS AND HOUSEHOLD SCALE SCORES

The food security scale was developed using statistical methods based on the Rasch measurement model. The model assumes an underlying continuum on which both items and households can be located. It assumes that the probability of a household affirming a specific item depends on the relative severity of the food insecurity of the household and that described by the item. For the food security scale, Rasch-based methods are used to determine item calibrations and household scores. An item's calibration represents the point on the scale at which there is a 50 percent probability that any given household at that severity level will affirm the item. Households with higher values on the scale than a particular item's calibration score have more than a 50 percent probability of affirming that item and conversely, those with lower values have less than a 50 percent probability of affirming the item. Item calibrations are calculated based on overall response patterns of all respondents. They are then used to calculate the severity score of households, based on the household's responses to the entire set of items.

There are two Rasch modeling approaches that have been used to calculate item calibrations, marginal maximum likelihood estimation (MML) and joint or unconditional, maximum likelihood estimation (JML). Details on both methods are presented below. Both methods have been used in various phases of developing the food security scale and assessing the consistency of the data with the statistical assumptions underlying the scale. JML methods were used to initially develop the scale and analyze data from the 1995 survey. Similarly, JML was used by the Economic Research Service for all their analyses and for the development of item and household scores published in the *Guide to Measuring Household Food Security, Revised 2000* (Bickel, G. M. Nord, C. Price, et.al, 2000). It served as the basis for the household scores in the 1998 and 1999 public-use data file as well. The MML method was used to reanalyze 1995 data, to analyze 1996 and 1997 data and as the basis for the household scores included in the 1996 and 1997 public-use data files. Here we compare the two methods as applied to the 1998 and 1999 food security supplement data.

Although the two methods are similar, the results produced are not identical. After a general description of Rasch models and a presentation of the results of the 1998 and 1999 estimation process of the food security scale, JML and MML methods are compared in detail and potential reasons for their differing results are explored. The central questions are whether the estimation method used makes any difference in the item calibrations that are obtained and, if so, which method is more appropriate for analyzing food security items and calculating household insecurity scores.

A. Using Rasch Modeling to Measure Food Insecurity

Rasch modeling relies on the assumption that the phenomenon being measured is continuous and can be portrayed as an interval measure. That is, the relative size of the intervals between household severity scores is meaningful, although the zero point is not. It assumes that each household has a score on a latent (unobserved) property that exists on a unidimensional scale. The model further assumes that each item that is used to form the scale is sensitive at a unique level of severity of food insecurity on this same unidimensional continuum. The probability of an affirmative response to any item is a function only of the respondent's level of food insecurity and the item's level of severity. It is assumed that the probability does not depend upon any of the other test items.

The distances between item scores and the ordering of items are meaningful in relative, but not absolute, terms. In other words, Rasch calibrations for a set of items are invariant relative to each other up to a linear transformation. Thus, comparisons of household scores or item calibrations require that the scales both be set to the same zero point. To accomplish this, the scales are adjusted so that the mean of the item calibrations is the same in both scales. The metric used by USDA's food security measurement project is based on a mean item calibration of 7. The size of the interval on the scale also can be made constant, for comparison purposes. The constant (called a scale factor or slope) used in the food security measurement project is 1. For some comparisons, scales may be adjusted so that the standard deviations of items are the same in both scales.

It is important to note that Rasch models do not assign scale scores to respondents with "extreme" response patterns. That is, if a respondent has no affirmative responses, we only know that the respondent's score is below the range measured by these items. Similarly, if the respondent says "yes" to all items, we know that the respondent's score is above the range that can be measured. Thus, in the food security application, households that answered no to all items (raw score of zero) did not have a scale score derived. Neither did households with severe food insecurity who answered yes to all items (raw score of 18). While we understand that the former group is food secure, we do not know how much more secure they are than households that answered only one item positively.

B. Food Security Item Calibrations, 1998 and 1999

Table 5.1 presents the item calibrations calculated for 1998 and 1999 using MML methods implemented by BILOG software[4]. The similarity in relative item severity between the two years as well as to earlier years is evidence of the stability of the measurement construct, and justifies comparison of prevalence rates across years. While in general the magnitudes of the calibrations are ordered similarly, there were two item reversals between the 1998 and the 1999 calibrations. In 1998, the survey question number 8a "adult cut or skipped meals, 3 or more months" had a calibration of 6.482 and question number 7 "children not eating enough" had a calibration of 6.738. By contrast, in 1999, while question number 8a "adult cut or skipped meals, 3 or more months" had an item calibration of 6.497, question number 7 "children not eating enough" had an item calibration of 6.383. This means that in 1998 "children not eating enough" was more severe than "adult cut or skipped meals, 3 or more months", but that the reverse was true in 1999.

There are a number of reasons why this might have happened. Sampling error might be responsible for the change in ordering but given that the differences are on the

[4] A detailed description of procedures used to calculate item calibrations and household scores using BILOG is available from ERS.

TABLE 5.1: 1998 and 1999 Food Security Item Calibration Values [a] (Discrimination (Slope) Parameter Set to 1.0)[b]

Survey Question Number [c]	Item Description	1998 Item Calibration	1998 Standard Error[d]	1999 Item Calibration	1999 Standard Error[d]
2	Worried food would run out	2.14	0.043	2.03	0.046
3	Food bought didn't last	3.40	0.037	3.10	0.041
5	Relied on a few kinds of low-cost foods for children	3.82	0.039	3.67	0.042
4	Couldn't afford to eat balanced meals	4.18	0.047	4.05	0.052
6	Couldn't feed the children a balanced meal	5.36	0.054	5.22	0.061
8	Adult cut size of meals or skipped meals	5.62	0.062	5.71	0.068
9	Respondent ate less than felt they should	5.75	0.049	5.87	0.055
8a	Adult cut or skipped meals, 3 or more months	6.48	0.055	6.50	0.063
7	Children not eating enough	6.74	0.042	6.38	0.048
10	Adult hungry but didn't eat	7.40	0.054	7.45	0.062
11	Respondent lost weight	8.29	0.062	8.41	0.072
13	Cut size of child's meals	8.57	0.088	8.89	0.103
12	Adult did not eat for whole day	8.73	0.113	8.74	0.132
15	Child hungry but couldn't afford more food	8.96	0.096	9.16	0.115
12a	Adult did not eat for whole day, 3 or more months	9.44	0.173	9.46	0.292
14	Child skipped meal	9.56	0.218	9.76	0.328
14a	Child skipped meals, 3 or more months	10.17	0.116	10.04	0.136
16	Child did not eat for whole day	11.33	0.230	11.55	0.281
	MEAN	7.00		7.00	
	STANDARD DEVIATION	2.584		2.695	

[a] Based on August 1998 Current Population Survey Food Security Supplement data.
[b] Adjusted to discrimination parameter (slope) = 1 and mean item score = 7.
[c] Items are ordered by severity reflected by 1998 item calibrations, which is slightly different from the order in the survey and that of 1999 calibrations. Question numbers are those in *Bickel, et al, (2000)* to facilitate comparison across years.
[d] Standard errors do not take account of possible inter-respondent correlation due to cluster sampling.

order of two standard errors, this is not very likely. The second possible explanation is that there is something in either BILOG software, or in the default settings when using that software, that renders item calibrations for BILOG less stable than those for the JML application. (This reversal is not seen in the JML estimates presented later in this section). The differences between the two programs are explored in more detail below. Third, a real change in the perceived meaning of items or their relationship to food security may have occurred.

The second such reversal occurred between “cut size of child’s meals” (question number 13) and “adult did not eat for whole day” (question number 12). In this case, however, the differences are within about one standard error and the reversal may well be the result of sampling variation.

Table 5.2 presents item calibrations based on separate estimations for 1995-1999 calculated using MML methods implemented by BILOG software. Each set of item calibrations has a mean set at 7 with the slope of the item characteristics curves at their inflection points set at 1. Each year’s scores are presented in the severity order reflected by 1998 calibrations. Given that the questionnaire underwent considerable reorganization in 1998, comparisons to that year’s data allow for an assessment of stability not only across years, but also across the questionnaire reorganization.

While the calibrations for all years are relatively similar in magnitude for most items and each scale spans between 8.7 to 9.5 units, there are some changes in the severity of items and differences in the order of severity when each year’s data are compared to 1998 and 1999. While the severity of the first two child items (question 5: “relied on a few kinds of low-cost foods for children” and question 6: “couldn’t feed the children a balanced meal”) remained approximately the same between 1995 and 1997, there was an apparent decline in their severity in 1998 and 1999. This is the result of the greater dispersion of items (higher standard deviation) in 1998 and 1999, reflecting a higher item discrimination. The higher item discrimination resulted from the introduction

TABLE 5.2: Comparison of Item Calibrations Estimated from April 1995, September 1996, April 1997, August 1998 and April 1999 CPS Food Security Data (Discrimination (Slope) Parameter Set to 1.0)[a]

Survey Question Number [b]	Item Description	1995 Item Calibration	1996 Item Calibration	1997 Item Calibration	1998 Item Calibration	1999 Item Calibration
2	Worried food would run out	2.55	2.47	2.37	2.14	2.03
3	Food bought didn't last	3.77	3.74	3.73	3.40	3.10
5	Relied on a few kinds of low-cost foods for children	4.34	4.37	4.53	3.82	3.67
4	Couldn't afford to eat balanced meals	4.02	3.99	4.10	4.18	4.05
6	Couldn't feed the children a balanced meal	5.61	5.62	5.82	5.36	5.22
8	Adult cut size of meals or skipped meals	5.54	5.47	5.56	5.62	5.71
9	Respondent ate less than felt they should	5.63	5.59	5.65	5.75	5.87
8a	Adult cut or skipped meals, 3 or more months	6.35	6.39	6.45	6.48	6.50
7	Children not eating enough	6.95	6.98	6.91	6.74	6.38
10	Adult hungry but didn't eat	7.21	7.20	7.29	7.40	7.45
11	Respondent lost weight	8.28	8.13	8.16	8.29	8.41
13	Cut size of child's meals	8.37	8.57	8.56	8.57	8.89
12	Adult did not eat for whole day	8.47	8.46	8.51	8.73	8.74
15	Child hungry but couldn't afford more food	8.63	8.98	8.95	8.96	9.16
12a	Adult did not eat for whole day, 3 or more months	9.02	9.05	9.01	9.44	9.46
14	Child skipped meal	9.65	9.67	9.29	9.56	9.76
14a	Child skipped meals, 3 or more months	10.19	10.15	9.86	10.17	10.04
16	Child did not eat for whole day	11.44	11.15	11.27	11.33	11.55
	MEAN	7.00	7.00	7.00	7.00	7.00
	STANDARD DEVIATION	2.449	2.450	2.40	2.584	2.695

[a] Adjusted to discrimination parameter (slope) = 1 and mean item score = 7.
[b] Items are ordered by severity reflected by 1998 item calibrations, which is slightly different from the order in the survey and that of 1995,1996,1997 and 1999 calibrations. Question numbers are those in *Bickel, et.al, (2000)* to facilitate comparison across years.

of internal screens and reordering of items in 1998, both of which tend to increase the consistency of response with item severity[5].

As the estimated severity level of these two items decrease, there is a resulting inversion in the severity ranking of the items following them. That is, in 1995, 1996 and 1997, question 5 ("relied on a few kinds of low-cost foods for children") was ranked as being more severe than question 4 ("couldn't afford to eat a balanced meal"). In 1998 and 1999 this ordering is reversed. In 1995-1997 the item calibrations for "couldn't afford to eat balanced meals" were lower than "relied on a few kinds of low-cost foods for children" by 0.326, 0.389 and 0.427 units respectively. In 1998 and 1999 the order of severity is reversed with item calibrations for "relied on a few kinds of low-cost foods for children" being lower than "couldn't afford to eat balanced meals" by 0.360 and 0.372 units respectively.

A similar pattern is observed between questions 6 and 8. In 1995-1997, item calibrations for question 6 "couldn't feed the children a balanced meal" were higher than those for question 8, "adult cut the size of meals or skipped meals". The opposite occurred in 1998 and 1999. Unlike the reversal described above however, the magnitude of the difference increased consistently through the years from .072 in 1995 to .494 in 1999. In 1996 and 1997 question 6 also fell below question 9 ("respondent ate less than they felt they should") in severity order.

It is most likely that these changes in relative item severity are the result of changes in the order of administration of the questions in 1998 and 1999. An exploration of the effects of screening on item calibrations (not shown here) has revealed that the internal screens increase the dispersion (standard deviation) of items by approximately three percent, but the effects of screening (initial, common, or internal) on relative severity of items are negligible and do not contribute in any substantial way to the

[5] Technically, the internal screens introduce inter-item dependencies that violate Rasch model assumptions. However, analysis of data from 1995-1997, which was not affected by internal screening, indicates that screening has negligible effects on relative item severities.

reversals of interest. Overall, the similarity of most item calibrations across the years and the consistency in their order of severity is indicative of the stability of the measurement construct across the years and across the significant reorganization of the questionnaire in 1998.

C. Comparing MML to JML Procedures

Item calibrations produced by BILOG's MML procedures are very similar to those produced by JML procedures[6] but are not identical. They differ somewhat because the conditions that the two methods impose on the maximum likelihood solution differ. Further, characteristics of the BILOG software in conjunction with characteristics of the food security data limit the precision of the MML estimates and require special handling. These differences do not threaten the meaning or reliability of the measure. The following compares the item calibrations developed by the two procedures and explores factors contributing to these differences.

To facilitate comparisons of item scores obtained using BILOG (MML) and JML procedures, item discrimination (or item slopes) must first be adjusted so that the dispersion of item scores from the two procedures are equal. With item slopes set to 1.0, as in Tables 5.1 and 5.2, the dispersion (as measured by standard deviation) of the BILOG item calibrations for the 18 items is less than that of JML calibrations. For 1998, the BILOG standard deviation is 2.584 and the JML standard deviation is 2.996, while for 1999 the standard deviations are 2.695 and 3.145 respectively.

The larger standard deviation of the JML estimates is consistent with a known upward bias on the dispersion of JML item calibrations. JML item calibration estimates are not statistically consistent. That is, as the sample size increases without limit, the JML item calibrations do not converge to their expected values. They are biased toward greater dispersion than the "true" item calibrations. However, the direction and approximate size of the bias is known. It has been shown that the JML inconsistency bias

---

[6] Joint maximum likelihood (JML) calibrations were calculated using an ERS adaptation of WINSTEPS (Lineacre and Wright, 1998). The ERS adaptation allows use of household case weights for estimating item parameters.

requires multiplication by a correction factor, (L-1)/L, to approximate consistency where L is the number of items (Andrich, 1988). Notice, however, that this would account for only about half of the difference observed in this case. Andrich's correction would be 17/18 (about -6 percent), or for households without children it would be 9/10 (-9 percent), whereas the observed differences in standard deviations are about 16 percent.

Because of the difference in dispersions of item parameters as estimated by BILOG and the JML method, we adjust to an equal standard deviation prior to an item-by-item comparison of severity[7]. Table 5.4 presents a comparison of the item calibrations calculated for 1998 and 1999 by JML procedures with the adjusted BILOG calibrations. While some of the differences between the two estimates are small, some are modestly substantial. For example, in the 1998 calibrations, four items have a difference between 0.10 and 0.16 including: "adult did not eat for a whole day for 3 or more months," "worried food would run out," "respondent lost weight," and "adult did not east for a whole day." In 1999 these four items continue to have a modestly substantial difference between the calibrations (0.13-0.17), along with two additional items which have a difference of –0.17 ("adult hungry but did not eat" and "child did not eat for whole day"). For both years the differences are greatest for the extreme items, either those that are most or least severe.

With a single exception, the ordering of the items is invariant between the two methods within each year. For the 1999 estimates, the JML procedures give the calibration of "adult did not eat for whole day" as 9.19 and the calibration of "cut size of child's meals" as 9.18, while BILOG gives the adjusted item calibration for "adult did not eat for whole day" as 9.02 and the calibration of "cut size of child's meals" as 9.19.

---

[7]BILOG scores were adjusted by a linear transformation to result in the same mean and standard deviations as the JML scores.

TABLE 5.4 Difference Between BILOG and JML Item Calibrations [a]

Item Description [b]	Item Calibration		JML – BILOG Difference 1998	Item Calibration		JML - BILOG Difference 1999
	JML 1998	BILOG 1998		JML 1999	BILOG 1999	
Worried food would run out	1.49	1.37	0.11	1.31	1.18	0.13
Food bought didn't last	2.79	2.83	-0.04	2.41	2.43	-0.02
Relied on a few kinds of low-cost foods for children	3.27	3.32	-0.05	3.03	3.10	-0.07
Couldn't afford to eat balanced meals	3.67	3.74	-0.07	3.48	3.54	-0.06
Couldn't feed the children a balanced meal	5.04	5.10	-0.06	4.83	4.90	-0.07
Adult cut size of meals or skipped meals	5.37	5.41	-0.03	6.22	6.26	-0.04
Respondent ate less than felt they should	5.53	5.56	-0.03	5.68	5.67	0.01
Adult cut or skipped meals, 3 or more months	6.42	6.41	0.02	6.46	6.40	0.07
Children not eating enough	6.66	6.70	-0.04	6.22	6.26	-0.04
Adult hungry but didn't eat	7.55	7.47	0.08	7.65	7.83	-0.17
Respondent lost weight	8.61	8.51	0.11	8.79	8.62	0.17
Cut size of child's meals	8.79	8.82	-0.03	9.18	9.19	-0.01
Adult did not eat for whole day	9.12	8.96	0.16	9.19	9.02	0.17
Child hungry but couldn't afford more food	9.24	9.27	-0.03	9.49	9.50	-0.01
Adult did not eat for whole day, 3 or more months	9.93	9.83	0.10	10.01	9.86	0.15
Child skipped meal	9.94	9.97	-0.04	10.17	10.20	-0.03
Child skipped meals, 3 or more months	10.63	10.68	-0.05	10.49	10.53	-0.04
Child did not eat for whole day	11.94	12.03	-0.09	12.12	12.29	-0.17
MEAN	7.00	7.00		7.00	7.00	
Standard Deviation (SD)	2.996	2.996		3.145	3.145	
Original SD ratio	1.16				1.17	

[a] Based on August 1998 and April 1999 Current Population Survey Food Security Supplement data. JML metric is in logistic units. BILOG metric is adjusted so that the standard deviation of the BILOG item scores is the same as that of the JML item scores.

[b] Items are ordered by severity reflected by 1998 item calibrations, which is slightly different from the order in the survey and that of 1999 calibrations.

Differences in the ways in which BILOG and JML applications employ maximum likelihood techniques cause these small differences in calibration results. The following explores some specific differences in program characteristics that may explain why the results vary.

Conceptual Differences

Conceptual differences between MML and JML procedures provide some explanation as to the variation in results. MML methods estimate the probability that a person with a particular score obtains a positive answer to a particular item with the individual person parameters conditioned away (or integrated out) (Andrich, 1988). This is based on an assumption about the distribution of the severity in the population. It is only dependent upon the values of all the item parameters for estimation and the assumptions about the distribution in the population. This means that the MML solution assumes that respondents are drawn randomly from a population of severities that is either a normal distribution or an arbitrary smooth distribution specified by the user (BILOG, 1990).

The JML procedure estimates both household and item parameters simultaneously. No assumption is required about distribution of the trait in the population. Household and item scores are estimated so as to maximize the likelihood of getting the observed matrix of responses under Rasch assumptions. The process begins with rough estimates of scores for each respondent and each item. These are then refined through an interactive Newton-Raphson process toward maximum likelihood until further parameter adjustments are smaller than specified convergence criteria.

Andrich comments that while the unconditional JML procedure is more efficient and converges faster than the MML procedure used by BILOG, it produces inconsistent estimates (in a statistical sense) for all of those estimates that are based upon a fixed number of items (questions). Statistical inconsistency is, ideally at least, undesirable in a measure because this means that even if the item and sample size is increased without limit, the estimate remains biased and does not converge to the value of the population parameter. This bias increases the dispersion of item scores estimated by JML methods relative to the dispersion of item scores estimated by MML methods. This is exactly what was found in the JML - BILOG comparisons described above. The size of this bias is known approximately, and is not generally problematic for the food security scale. It should, however, be kept in mind when comparing results from scales based on a subset of the items, such as the six-item standard short module.

Problems Introduced by MML Assumption of Smooth Distribution

As estimated under MML assumptions, the severities of the child item scores (or some of them, at least) are distorted relative to those of the adult and household item scores because of the assumption inherent in the MML method of smooth distribution of severity in the population. The distortion arises because in fact, two populations—households with and without children—are intermingled in the food security data. This wouldn't matter if all households got the same set of items, but because all of the child items are estimated based only on the population of households with children, the two sets of household score groups are interspersed, but represent the two populations. (And, the more severe household score groups are all based only on households with children.) Depending on the exact method used to assess the distribution, BILOG may adjust child items disproportionately relative to other items to try to smooth the distribution. BILOG (MML) and JML procedures produce item calibrations that are almost perfectly linear with respect to each other if the universe is restricted to households with children or to households without children. However, when the population is mixed (households with and without children), the calibrations are no longer perfectly linear. Thus, the BILOG (or generic MML) assumptions of a smooth distribution (inappropriate when the two household types are mixed) account for most of the non-linearity between BILOG and JML item scores described above. The MML assumption regarding distribution of the measured trait in the population may also be violated by the screening procedures in the Food Security Supplement. This violation may also contribute to the difference between JML and MML scores. Even so, the differences are not very large and in most cases have little practical implication. But they are large enough to be noticeable.

Convergence

The major problem with using BILOG to estimate item calibrations from the food security data is lack of convergence in the Newton Raphson phase. The non-convergence problem is a result of the dependent frequency follow-up items. Those items and their base items oscillate in opposite directions in alternate steps of the iteration process. The size of the oscillation is quite large, in the range of 2 to 4 logistic units. If the frequency follow-ups are removed, then the calibrations will converge to 0.001 in 5 to 15 NR steps. Similarly, if the base items of the frequency follow-ups are removed and their frequency follow-ups included, the calibrations will converge.

In an attempt to resolve this problem, the free option (an option on the Calibration command) was tested. It was thought that use of an empirically derived prior distribution as opposed to a program-imposed normal distribution might resolve the convergence problem. Although the results are not shown here, these calibrations also failed to converge on a stable estimate. To address the issue of non-convergence, the two item sets were estimated separately (with frequency follow-ups excluded in one run and their base items excluded in the second), and then the metrics were equated by the common items. Interestingly, the set of scores derived by this process is very nearly linear with the scores at the end of the so-called EM estimation phase (after 10 cycles or so). Thus, the item dependencies appear to be problematic only in the Newton-Raphson phase, not in the EM phase.

It is also worth pointing out that the BILOG software and manuals are oriented to item response theory (IRT) users in general and not to Rasch modelers in particular. This makes it more difficult for the novice user to use BILOG to estimate a one-parameter logistic (Rasch) model than a JML program such as WINSTEPS, which is specifically oriented to Rasch's own views of measurement rather than IRT more generally.

D. Households Scores, 1998 and 1999

Household scores on the food security scale are also calculated based on the Rasch model. One result of the Rasch assumption that items discriminate equally is that, for a given set of items, the household's scale score depends only on the number of items affirmed, not on which items are affirmed. Thus, for households with no missing data, households with the same raw score will be assigned the same scale score. If respondents are not all given the same set of questions, the scale scores depend on the severity (as indicated by the item calibration) of the questions that the respondent answers, as well as the number of items affirmed. The food security supplement includes 18 questions for household with children and 10 questions for those without. The Rasch model takes these differences into account, assigning scores to both types of households that are comparable even though they responded to different subsets of questions. The model also adjusts the household scale scores for households that failed to respond to one of more of the applicable questions.

Once item calibrations have been determined, household scale scores for households with no missing values can be calculated for each raw score. Identical household scores are produced by BILOG and JML software provided the same item calibrations are specified. After reviewing all item calibrations, it was determined that the JML calibrations would be used as the "standard". Although JML estimates are somewhat biased toward greater dispersion than their true value, ordinality of items is preserved and relative severities (proportional size of intervals between items) appear to be either unaffected or minimally affected by the statistical inconsistency. These latter two important characteristics cannot be assured in MML estimates except through a complex process of fitting multiple models. Table 5.5 presents the household scores that were used for both 1998 and 1999 public-use data files, which were based on JML methods and the 1998 data.

Table 5.5 Household Food Security Scale Scores, 1998 and 1999

Number of "yes" responses		Household Scale Scores 1998 and 1999	Food Security Status Category
Household with child	Household without child		
1		1.428	Food Secure
	1	1.723	
2		2.560	
	2	3.101	
3		3.405	Food Insecure without Hunger
4		4.138	
	3	4.232	
5		4.138	
	4	5.234	
6		5.430	
7		6.024	
	5	6.155	
8		6.606	Food Insecure with Hunger
	6	7.068	
9		7.179	
10		7.738	
	7	8.002	
11		8.276	
12		8.794	
	8	8.976	
13		9.306	
14		9.837	
	9	10.149	
15		10.423	
16		11.133	
17		12.157	